

Internet Engineering Task Force
Internet-Draft
Updates: [4456](#) (if approved)
Intended status: Standards Track
Expires: April 21, 2013

J. Scudder
Juniper Networks
October 18, 2012

Considerations for Route Reflection and EBG
draft-scudder-idr-ebgp-rr-02

Abstract

Although originally conceived of as a purely IBGP device, in some cases a route reflector may function as an EBG speaker in addition to its role as envisioned in [RFC 4456](#). When it does so, just like any other EBG speaker it must advertise its routes to its IBGP peers. This document updates [RFC 4456](#) by explaining what behavior is required of a route reflector that also functions as an EBG speaker. It also clarifies the use of the ORIGINATOR_ID and CLUSTER_LIST attributes in this environment.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 21, 2013.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect

Internet-Draft

RR and EBG

October 2012

to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

1. Introduction

Although originally conceived of as a purely IBGP device, in some cases a BGP [[RFC4271](#)] route reflector may function as an EBG speaker in addition to its role as envisioned in [[RFC4456](#)]. When it does so, just like any other EBG speaker it must advertise its routes to its IBGP peers. This document updates [RFC 4456](#) by explaining what behavior is required of a route reflector that also functions as an EBG speaker. It also clarifies the use of the ORIGINATOR_ID and CLUSTER_LIST attributes in this environment.

The cardinal observation is that the functions outlined in [[RFC4456](#)] apply exclusively to "reflected" routes, that is, IBGP routes that are propagated to IBGP peers.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

2. Terminology

In addition to the terms defined in [[RFC4271](#)] [Section 1.1](#) and in [[RFC4456](#)], this document makes use of the following:

ASBR: Autonomous System Border Router. See EBG Speaker.

RR: Route Reflector.

Redundant Route Reflector (or Redundant RR): Another route reflector in the same cluster as the route reflector under consideration, when both route reflectors are configured with the same CLUSTER_ID.

EBG Speaker: A BGP speaker that has one or more EBG peerings,

and thereby learns one or more EBGp routes. (If no routes are learned it is still an EBGp Speaker, but this is a case of "a tree falling in a forest with no one to hear it.") ASBRs are EBGp speakers, although not all EBGp speakers are ASBRs.

3. Updates to [RFC 4456](#)

When deciding how a route reflector that is also an EBGp speaker should propagate EBGp routes into IBGP, the key observation is that [\[RFC4456\]](#) deals only with "reflected" routes, i.e. IBGP routes that are propagated into IBGP. For EBGp-learned routes, the BGP speaker is the only source of routes for its AS. For this reason, the restrictions and assumptions that apply to reflected routes do not apply to EBGp-learned routes. For the purposes of such routes, the BGP speaker functions as a normal IBGP router. For example, the [\[RFC4456\]](#) stricture against modifying the NEXT_HOP, AS_PATH, LOCAL_PREF, and MED attributes does not apply to EBGp-learned routes that are propagated into IBGP.

Specific updates to [\[RFC4456\]](#) are:

- o The speaker MUST NOT add a CLUSTER_LIST to EBGp-learned routes when advertising them into IBGP.
- o The attributes ORIGINATOR_ID and CLUSTER_LIST MUST NOT be sent to EBGp peers. If received from an EBGp peer, these attributes MUST be ignored and not propagated further; an error message MAY be logged.

4. Deployment Considerations

If route reflectors are deployed in an Autonomous System such that no two route reflectors have the same CLUSTER_ID, then there are no "redundant route reflectors" (as the term is used herein) and thus, the considerations regarding redundant RRs below are moot.

A RR that serves as an EBGp speaker SHOULD have an IBGP peering with any redundant RR. It SHOULD advertise the same EBGp-learned routes over this peering that it advertises to any other IBGP peer. It MAY suppress reflection of any IBGP-learned routes to the redundant RR.

(Recall that according to [\[RFC4456\] Section 8](#), such routes would be ignored by the redundant RR anyway due to a loop in the CLUSTER_LIST.) The peering MAY be omitted if the route reflectors in question are control plane-only devices not in the forwarding path of any traffic, or if the network in question uses some form of tunneled or label-switched forwarding. The cost of omitting the peering is that certain low-probability failure modes may cause unnecessary unreachability -- specifically, if the EBGp-speaking RR were to lose its session to one or more of its RR clients, reachability to the EBGp-learned routes would be preserved if a session remained up to its redundant RR peer. (Similar considerations apply even to route reflectors which do not have a colocated EBGp speaker function, but

such are beyond the scope of this document.)

[5.](#) IANA Considerations

This document makes no request of IANA.

[6.](#) Security Considerations

This clarification to BGP does not change the underlying security issues.

[7.](#) Acknowledgements

The author would like to thank Serpil Bayraktar, Jeff Haas, Senad Palislamovic, Yakov Rekhter, Jim Uttaro and Kaliraj Vairavakkalai for their input.

[8.](#) Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

[RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), January 2006.

[RFC4456] Bates, T., Chen, E., and R. Chandra, "BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)", [RFC 4456](#), April 2006.

Author's Address

John Scudder
Juniper Networks

Email: jgs@juniper.net